[SQUEAKING]

[RUSTLING]

[CLICKING]

JONATHAN GRUBER: All right, let's continue our discussion of factor markets. Remember, to set the stage, we are now in the second half of the course, talking about the various elements that underlie the supply and demand model that we developed in the first half of the course.

And the first question we're tackling is, where do w and r come from? Where do these prices that are the prices of firm inputs, where do they come from? And they come, of course, from markets, markets of demand and supply, markets' four factors.

So let's start by doing the labor market equilibrium. Let's start by doing the labor market equilibrium. And then we'll talk about what effect a minimum wage has in that equilibrium. So let's look at figure 16-1. Figure 16-1 puts together the labor market equilibrium.

We have a demand curve. The demand curve comes from firms because in factor markets, firms are now the demanders, and people are the suppliers. So we flipped the tables from a typical market. In a typical market, in a consumer goods market-- I should state it in order-- a consumer goods market, firms drive the supply curve and consumers drive the demand curve. That's no longer true.

Now firms drive the demand curve and consumers drive the supply curve. So the demand curve comes from the marginal product of labor. So we have a demand curve in labor market equilibrium. Here, we're saying the demand curve is 70 minus w. We're saying that as the wage goes up, firms want fewer workers. And we can see that from our earlier firm optimization.

But there's also a supply curve. Now, we talked about last time, the shape of the supply curve here is indeterminate. And that's because income effects offset substitution effects when leisure is a normal good. That's the key lesson from last lecture, is that when leisure is a normal good, income effects offset substitution effects. So we don't know the shape of the supply curve.

That said, we're going to pretty much, for the rest of the semester, assume upward sloping labor-supply curves. We're going to pretty much assume substitution effects offset income effects because it's a hell of a lot easier to work with. It's a mess when they both slope down. So pretty much going to assume that substitution effects dominate, and the labor supply curve is upward sloping.

So here, for example, we have a labor-supply curve of the form that labor equals w minus 10, so a higher wage leads to more labor supply. So a higher wage leads to more labor supply. This should be l equals 70. So this is the demand curve. The demand curve is l equals 70 minus w. The supply curve is l equals w minus 10.

So to solve for equilibrium, we simply set those equal. We know how to do that now. And you find the equilibrium is at 30 units of labor with a wage of 40. That's the point at which firms are happy to hire workers, and workers are happy to work. So it's the same kind of equilibrium of goods markets. Now, it's just in a factor market.

Now, let's ask what happens when we add a minimum wage? What happens we add a minimum wage, which is a standard labor market intervention all around the world. Some of you may have even enjoyed minimum wage jobs at some point in your past lives, or hopefully not in your future lives but at some point in your past lives.

It's a common feature of the developed world to have minimum wages. way the minimum wage work, is the government just sets a wage floor below which employers cannot pay their employees. They set a wage floor.

So for example, imagine the government came in and set a wage floor of 50, set a minimum wage of 50 above the market equilibrium wage. Sp the new equilibrium now, is the new outcome, the new market outcome, is actually at 0.2. Why? Well, because that higher wage, firms want to hire fewer workers. And it doesn't matter what workers want. There's no jobs for them.

So firms say, fine, at that higher wage, I only want 20 workers. Workers say, well, gee, at that higher wage, I want to work, actually, I want to work much more than that. OK, Andrew, in the future diagram, we should have the x-intercept for the excess supply. They want to work much more than that, so you end up with an excess supply of labor.

Now when we talked about supply and demand equilibrium, we talked about how disequilibrium outcomes resolve themselves through market adjustments. Well here, this disequilibrium outcome cannot resolve itself. The market can't adjust. The market's stuck because of the regulation. So we end up in a situation, we're out of equilibrium. We have excess supply.

In layman's terms, we call that unemployment. We end up with unemployment. We end up with people who want to work but can't work. And that causes an inefficiency or deadweight loss.

You see, before the policy, worker surplus was-- or, I'm sorry, firm surplus. I'm sorry. Before the policy, firm surplus-- because firms are the consumers here --was a plus b plus c. That was the consumer surplus, or in this case, consumers are firms. So really call that more the demander surplus.

The supplier surplus was d plus e. With the minimum wage, we've now transferred the area b from firms to workers. So the firm surplus, the demander surplus, has fallen all the way to a. So we've transferred the area b. But we've created, in doing so, a deadweight loss of c plus e . Standard kind of welfare analysis we've done so several times before now.

So what we've done with this policy, is we've transferred the area b from the demanders who are firms, to the suppliers who are workers. But in doing so, we've created a deadweight loss C plus E. Why? Because those are trades that would be efficient to make. They're not being made.

There are workers that would happily work at a wage below 50, and firms would be happy to hire them at a wage below 50. But we're not letting that transaction happen. That's wasted economic efficiency or deadweight loss.

This is the pretty standard analysis of minimum wage. And this says that, well, look, when you put a minimum wage, you get a trade-off. You get a transfer to workers, which you may like, but you create a deadweight loss because you create unemployment.

But actually, starting about 30 years ago, research started finding that minimum wages didn't cause unemployment. Actually, some research studies suggest you put a minimum wage, employment goes up, but certainly, there was very little evidence that goes down. This was quite striking because how could that happen?

Now, one way it could happen, of course, could be if the demand curve was perfectly inelastic, then a minimum wage would not lower employment. You can see that in this graph. I haven't drawn this out, but you can see in 16-2, imagine taking that blue line and making it vertical, you'll see that the minimum wage wouldn't cause unemployment.

So one thing that could be going on is demand could be perfectly inelastic. But in fact, there's lots of studies which say that's not true. Demand for labor is not perfectly inelastic. So that can't be the explanation.

So we said an alternative explanation, which is that just like goods markets can be non-competitive, labor markets can also be non-competitive. And so we think the explanation is non-competitive labor markets.

So when we thought about firms selling goods, we said, look, in a competitive market, firms face perfectly elastic demand. Any given firm faces perfectly elastic demand. Well, likewise, when we think about a firm hiring workers, we think that if labor markets are competitive, any given firm faces a perfectly elastic supply curve. Let's think about the intuition.

The intuition in a perfectly competitive market is, look, if I'm selling tchotchkes next to the other guy selling tchotchkes, then I can't raise my price above or below his because if I raise it above, I'll sell zero. Below, I won't be able to meet demand.

Well, it's the same thing when hiring workers in a perfectly competitive market. Imagine a perfectly competitive labor market, which is fast food workers. There's a bunch of fast food restaurants near each other. If McDonald's pays $0.50 more than Burger King, everyone leaves Burger King and works at McDonald's. If Burger King pays $0.50 more than McDonald's, vice versa.

So you have essentially each firm faces a flat supply curve because workers can just move. It's an easy substitution between them. Just like it's easy to Substitute between statues of the Eiffel Tower, little mini statues, it's easy to substitute been working at McDonald's and Burger King. It's the same job, but they're near each other.

However, what we've realized is that may in fact, not be true. What if just like goods markets can be non-competitive, what if labor markets can be non-competitive? Remember the assumptions we required for a perfectly competitive goods market. We required perfect information and no transaction costs and all goods being the same.

Well, for a perfectly competitive labor market, the same requirements are in place. You need perfect information, you need no transaction costs, and all jobs to be the same. Well, that's not true. In particular, the first isn't true.

We don't have perfect information because jobs don't always post their wages, typically don't post their wages. And typically, you don't know what you can make at another firm. Once you don't know what you can make at another firm, once it's costly to learn that, that's a transaction cost that interferes. That's one form of transaction costs that can interfere with what we think is a fairly competitive market.

Now, let's go further. Let's think about the most extreme case, which was the old mining towns we used to have in the West, in the US. In the US, the way it worked is you discovered gold. You set up a town around the gold mine, and everyone in that town would work in the gold mine, or they would work selling stuff to be worked in the gold mine.

Well, in that case, you definitely didn't face a competitive labor market because that was the only place to work. There was one employer, the gold mine, and if you wanted to work, you had to work for that one employer. So instead of the case of a monopoly, we had a case of a monopsony. A monopoly is when only one firm sells the good. A monopsony is when only one firm hires the workers.

Let's suppose further-- let's just look at this town as having a bunch of gold mine workers. Let's suppose further that just as we assume initially, that monopolies can't price discriminate, let's assume monopsonies can't wage discriminate. This is actually kind of a more reasonable assumption. It was always weird to say why couldn't monopolists charge different people different prices?

It's a more reasonable assumption to think firms can't pay similar workers different wages. They'd be pissed. If you're paying two workers doing the same job, different wages, they're going to get upset. People care about equity. It's more reasonable to think for the same job, you have to pay the same wage. I think that's a much more reasonable assumption than it is a monopoly case.

I mean, you can pay managers more than mine workers. But it's reasonable to think you have to pay all the mine workers the same thing. What that means is that just as when you're a monopoly, lowering your price might hurt your profits through the poisoning effect, when you're a monopsony, raising the wage will hurt your profits through an equivalent poisoning effect.

So let's think about the firm's decision. The firm wants to set the marginal cost of the next worker equal to the marginal benefit of the next worker. What is the marginal benefit? Well, the marginal benefit is still the marginal revenue product of labor. That hasn't changed. That's true regardless of the market. The next worker produces the next marginal revenue to produce the marginal product of labor.

What changes the marginal cost? Before the marginal cost, was the wage. In a perfectly competitive labor market, I didn't say last time I was assuming that, but I was. In a perfectly competitive labor market, the wage is the cost. But in a monopsony market, that's no longer true. And to see that, I want to say, what's the marginal cost? The marginal cost is the marginal expenditure to get another unit of labor.

Well, what's the expenditure to get another unit of labor? Your expenditure -- is what's your total wage bill? I'm sorry. What's your total wage bill? It's the wage you pay by the number of workers you have. But the number of workers you have is a function of the wage you pay, right? The more you pay, the more workers will want to work with you.

So if we differentiate this and get marginal expenditure, which is the firm's marginal cost of hiring another unit of labor, the marginal expenditure is w plus dw dl times l. In other words, it's this term, which is positive, the marginal expenditure is the wage, plus this second term, which is how much do you have to raise the wage to get another worker?

It's the flip side of the poisoning effect. The poisoning effect was, as a monopolist, if I want to sell one more unit, I face a trade-off. I sell the additional unit. That's good, But I have to lower my price to sell it. That's bad.

Here, if I want to hire one more unit of labor, there's an equivalent trade-off. I get one more worker. That's good. But I have to pay all my workers more to get him. That's bad. Why? Because I can't wage discriminate.

And it makes sense. If I want to expand the number of-- once again, I think it's much more intuitive than the monopoly case. If I want more workers, I've got to raise the wage. Why was this not an issue for the perfectly competitive firm? It was not an issue with the competitive firm because they faced a perfectly elastic firm labor-supply curve. Workers weren't different in McDonald's and Burger King, so they could just go back and forth. But a monopsonist doesn't face that anymore. They face the true market labor supply-curve, which is not perfectly elastic.

So to see this, let's go to figure 16-4. So 16-3, that's a perfectly competitive case. The marginal benefit of hiring another worker is the marginal revenue product of labor is 16-3, the marginal benefit the marginal revenue product of labor. The marginal cost is expressed by the wage, which is determined by a perfectly competitive market.

But 16-4 is a monopsony. Now we have the same marginal revenue product of labor, which we defined as 70 minus l. We used the same example did before, 70 minus l. Well, we're flipping this. So now we can write the marginal revenue product of labor curve, as simply w equals 70 minus l, so just inverting that. That's the firm's demand curve or marginal product of labor curve.

But what about its marginal expenditure curve? Well, if you look at the supply curve, the supply curve was w equals l plus 10. So we can now say, the firm's expenditure function, labor is w-- I'm sorry. The firm's expenditure function, is l times l plus 10. That's its labor expenditures. Stare at that for a minute. The number of workers times the wage. The wage comes from this equation.

Differentiate that with respect to l, and what do you get? If you differentiate If with respect to l, you get 2 l plus 10. And that is the firm's marginal expenditure curve, which is above the supply curve, as you see in figure 16-4.

So what does the monopsonist do? The monopsonist sets-- goes to the point where marginal cost equals marginal benefits. What's the marginal cost? The marginal cost is the marginal expenditure curve. What's the marginal benefit? The marginal revenue product curve.

So it says, I want to hire 20 workers. Not the 30 I was hiring when the market was competitive but 20. What wage am I going to pay them? Well, just like the monopolist needs to respect the demand curve, the monopsonist needs to respect the supply curve. So a wage you can pay them, when you read off the supply curve, you're going to pay him 30 instead of 40.

So what's going on here? You're hiring fewer workers at a lower wage than the equilibrium would dictate. Why? Because if you want more workers, you got to raise the wage, and that's not profitable. So if you're at the point where you're paying 30 to get 20 workers, and from that point, you want to hire more workers, say, well, look, even at that point, I'm making money on the next worker.

But the problem is to hire that next worker, you have to pay higher wages to all your previous workers, and it's not worth it. This is a hard intuition, but it's the flip side of the poisoning intuition. So go back to your monopoly case and work it through. You'll see it's all just the opposite.

Now, well, I need to teach you the math and the graphs because that's what we test you on. The intuition here is actually pretty clear, I think, clearer than the monopoly case. Just think about the intuition for a second.

Think about me at MIT. There's many universities I could work for. OK, I don't mean to sound arrogant, but there's a lot of universities out there. Presumably some others would want to hire me. Yet, MIT has monopsony power over me. Why? What gives MIT monopsony power over me?

It's nothing nefarious. Think about typically any firm has monopoly power over their employees. Why? Well, yeah?

AUDIENCE: Because you don't know how much you'd get paid anywhere else.

JONATHAN GRUBER: I don't know how much I'd get paid anywhere else. That's one reason. So there's an information reason. What else? Now, actually let's compare that. Actually, I know what I'd get paid in California because California public schools have to make their wages public.

So let's say I could overcome that. Actually, if you go online, you want to raise pay at a public university in America, you can figure that out. OK, so let's say I could solve that problem. But there's other reasons why MIT has monopsony power. Let's say I'm thinking of moving to California. Yeah?

AUDIENCE: Well, if you're already established, [INAUDIBLE]?

JONATHAN GRUBER: Well, I'm established. Like, if it was-- I have to fucking move, right? If I want to go to California, I got to move. That sucks. I got to pack up the house. I got to leave all my friends. In particular, my wife might have a job that she likes here. My kids might like being in school here.

So there's a huge transaction cost. Transaction costs create monopsony power. It's not a mining town. MIT's not a mining town. But it still has market power over me because of transaction costs. There's also the fact I like teaching at MIT. I've got my lecture set. If I went to some other university, I'd have to dumb them down because the students aren't as smart as MIT. I don't want to do that.

My office has a window view. It's pretty awesome. My window has a river view. My office has a window view, Jesus. My office has a river view through the window. OK, that's pretty awesome. That'd be hard to replicate at a lot of other universities.

Basically, every job, changing jobs comes with transaction costs. Transaction costs create monopsony power. MIT does not have to pay me the market wage because of these transaction costs and because of information limitations.

So it's not crazy to think-- indeed, we now think that basically, monopsony is a better description of almost all labor markets than is perfect competition. This is a real sea change in the last 20 years is a real takeover of the monopsony model in labor economics. That we think it's probably much more likely that firms have monopsony power than that they're perfectly competitive. Yeah?

AUDIENCE: When you first figure it out, we don't have to know some of the other expenses or transaction costs or whatever. So it's still monopsony [INAUDIBLE]?

JONATHAN GRUBER: In the new-hire market, so that's a great point. There's still the information issue, which is I'd have to go and collect a bunch of offers. But other than that, you're right. And indeed, if you look at salaries of academics, new-hire salaries, if you look at that salary of academics, you can basically predict them pretty well based on essentially, when do they get their last offer?

So new-hire salaries are pretty flat, and they don't get raises when they get an offer from another university. So basically, MIT will underpay you until you get an offer from another university. And then they'll-- because that reveals information, and then MIT can then do that, OK, good question.

All right, now, what does this mean for the government? Well, this means that there is scope for gains from regulation. Because look, this monopsonist has created a deadweight loss. We also have 30 workers. So there's a deadweight loss triangle from only hiring 20 workers. That says there's scope for optimal regulation, just like there's scope for optimal regulation of monopolists.

And what would optimal regulation look like? Well, it could look like a minimum wage. So imagine in figure 16-5, we set a minimum wage equal to the market wage. Well, now think about the firm's decision. For any workers to the left of l comp, there's no poisoning effect because they can't lower the wage. So they might as well hire the competitive level.

What has the minimum wage just done? It has raised employment. So this is how you can get this weird result that a minimum wage can raise employment. Because what the minimum wage has done, has essentially eroded monopsony power or eroded the negative effects of monopsony by getting rid of this poisoning effect, by saying, well, firms might as well hire l comp. The reason they're hiring lower than l comp before was because they could lower the wage, but now they can't.

So that is how a minimum wage can raise employment. Now, of course, so in some sense, you can get a win-win government policy. You could transfer surplus to consumers and raise employment. Everybody wins. And that's how government policy can make things better in with market failures.

Now, of course, that isn't necessarily true. What would happen if you set the minimum wage-- well, let me ask first, what happened if you set the minimum wage too low? What if you set the minimum wage below w1? What would happen if the government put the minimum wage, and it was below w1? Yeah?

AUDIENCE: There would be no effect.

JONATHAN GRUBER: No effect, very important. The private market will always get around government regulations if it can. A government regulation that doesn't bind, it's no government regulation at all. It sounds obvious, but it's not. It's important to remember the first thing you have to ask of good regulations is, do they matter?

Then basically, as you think about any minimum wage between w1 and even past w comp, they will actually improve welfare. But as it gets high enough, the minimum wage could actually cause firms-- once when a wage gets above the level on the marginal revenue product curve that corresponds to l1-- so you see where the l1 dashed line hits the marginal revenue product curve. Once the wage gets above that level, then you're making things worse. You can end up hiring fewer workers than the monopsonist would hire.

So once again, the way to think of government regulation, once again, come take 1441, we'll spend all semester on this, but the way to think about it is, the market knows best, unless there's a market failure. The market failure raises the potential for the government to improve things but doesn't prove government will improve things. It does not prove government will improve things.

You need to actually study and examine the government policy to see if it makes it better or worse. For minimum wages, basically, the consensus now is, for minimum wages in the range that minimum wages are set in America, which is probably about $8 and $12 an hour, where most minimum wages are, the consensus seems to be that they are a good thing. That basically, they can cause benefits to workers without actually reducing employment.

Once you get above 12, the evidence gets more mixed. So we now are seeing a series of 15 hour wage dollar minimum wage, including Massachusetts. The evidence on those is a little more mixed. We need a little more time to study. But there's some level at which a minimum wage, definitionally, is high enough, it will cause unemployment. That doesn't mean minimum wage is bad. It just means at that point, there's a trade-off.

So if you go back to figure 16-2, 16-2 does not imply we should not have minimum wage. 16-2 just says there is a trade-off. You lose efficiency, but you redistribute to workers. In a standard model, that's bad because we rate consumer and producer surplus the same. But in the real world, it might not be. We might value consumer surplus more than producer surplus.

So the bottom line is, the minimum wage isn't-- in no situation is it unambiguously bad. It may actually be unambiguously good. And it depends on the structure of the labor market. OK, questions about that?

So once again, the idea would be MIT is making excess money off me because I don't want to move. If the government came in and raised my wages, not too much, that would make me better off, and it would just come out of MIT's profits. It would just be a transfer. But I'd be better off. And they'd maybe-- because they pay higher wages, they may hire another professor. And that would improve efficiency. All right, so with monopsonists, minimum wage is just induce efficient transfers.

Now, that's all I want to say about labor markets. I want to turn now to capital markets, which are a degree more complicated and therefore, a degree more interesting. The main reason they're complicated, is that capital is not tangible the way labor is. Labor, you feel it. You've done it. It's an hour of work. You know what work is.

Capital, what is it? It's machines. It's buildings. It's this vague thing. So we're going to now define capital more specifically by thinking about where capital comes from. Where does the money come from to buy machines, to finally to invest?

Well, it goes all the way back to our farming agricultural economy. Farmers, every year, face a decision. Are they going to consume all of their crop, or are they going to save some of it to plant to create next year's crop? So farmers face the decision of diverting today's consumption to tomorrow's consumption, or next year's consumption.

How much of today's consumption do they want to give up, how much less do they want to eat this year, to have more to eat next year? And that determines how many seeds they want to set aside. That's what capital is. Capital is essentially, what you take from today to have in the future.

When you buy a machine, you are spending money today to get the machine so that over the next few years, it produces stuff. That's no different than saving your seeds. That's what capital is. It's essentially the diversion of current consumption towards future consumption.

Now, unlike the farmer, firms today don't get the money necessarily from themselves like the farmers did. Firms get their capital in what's called a capital market. Well, I just wrote that there. Firms get their capital in the capital market. What is the capital market?

Well, that's a market where there's a demand for capital, which comes from how much-- which I'll derive in a few minutes, comes from how much firms value the capital and the supply of capital. Well, where does supply of capital come from? What is capital? It's basically the money you use to invest in the future. Where does it come from? It comes from us. The supply of capital is savings.

So the money firms use to buy their machines, they borrow from people. And that is how the capital market works. So when we talk about capital, well, really think of it fundamentally, as financial capital. Don't think of it anymore as machines and buildings. Think of it as financial capital. The money you want to invest to divert today's consumption to tomorrow's consumption, it comes from the savings that we do, which is our decision to divert today's consumption to tomorrow's consumption.

So just as workers decide how hard to work, and that drives the supply of labor, people decide how much to save. And that drives the supply of capital.

So let's see that in figure 16-6. Figure 16-6 is the capital market equilibrium. The demand for capital comes from the marginal revenue product of capital. Just like the demand for labor comes from the marginal revenue product of labor, the demand for capital comes to the marginal revenue product of capital.

How much will the next unit of capital, the next unit of financial investment, what will it yield in terms of the value of what I'll produce? How much will the next machine increase the quantity of goods and what can I sell those goods for? So the same logic that applies to labor, applies to capital.

The supply of capital comes from workers' saving. Why do workers save? They save because they can get a return to that saving in the form of interest. i is the rate of interest. That's the return that workers get on their savings.

So in other words, you could think of i, the technical definition of the interest rate, is the rate that firms have to pay you to get your money. i is the rate that firms have to pay savers to get them to save. So for example, if the interest rate is 10%, then if I want to borrow $1 from a household this period, I have to pay them back $1.10 next period.

By the way, one parenthetical comment, time is a very elastic concept in this class, man. Don't be uptight in your thinking, like the Dude says. Basically, time, sometimes I'll say today, tomorrow. Sometimes I'll say this year, next year. It's all the same, man. It just means period one versus period two.

So don't get too caught up in the words I use, just think of this as period one versus period two. That's the key. And I'll try to use it as much as I can. But if I say today, tomorrow, this year, next year it's all the same thing. So don't get too caught up on the terminology.

Now basically, the idea is that the demand curve is downward sloping. As the interest rate rises, firms want less capital. Because why? Because interest rate rising, means that the opportunity cost of employing machine has gone up. Because what can you do instead with your money? Hire a worker.

So a downward-sloping demand curve for capital, is basically the rate at which firms are going to trade-off capital for labor. As interest rate goes up, they'll say fewer machines, more labor. Interest rate goes down, they'll say less labor, more machines. The supply curve slopes up because basically, the more you pay people to save, the more they'll save.

Now, let me pause for a second. We have a whole series of courses, in course 15, about finance. I'm doing finance in like a lecture and a half here. And if you find this interesting, you take 15401, 15402, this whole series of courses. You take course 15. By the way, since all of you should really go be econ majors, you get economics elective credit for a couple of those courses. You can be any kind of major and take some course 15 courses if you want to learn more about finance.

But let's just talk about-- this weird thing I'm talking about, what does it mean that people loan money to a firm? Because we don't really necessarily loan money to a firm. In fact, firms get their money from people in three different ways. There's three different ways we can essentially transfer our money to firms.

One way, is we could put it in the bank, You put your money in the bank. The bank pays you an interest rate. Firms then go to the bank and borrow the money from the bank. So we're not directly giving the money to the firms. We're doing it through a financial intermediary, which is the bank.

So the most common way people get money to firms is by having money in the bank. You don't know that. You just think you've got a bank account yielding whatever, but that's what you're doing. You're just giving money to firms.

The second way you can do it, is you can actually directly loan the money to the firm by buying corporate bonds. A corporate bond is literally a loan to a firm. Now, typically, most of us who aren't super rich won't buy a corporate bond. We'll buy a financial instrument that holds corporate bonds.

So I have a Vanguard Bond Fund. Vanguard is a company that holds financial instruments. So I own a little piece of a lot of bonds. Bonds are millions of dollars. I'm not going to own a bond. But I own a little piece of a bunch of them. So that's the way you directly loan to a firm.

And then finally, you can buy stocks, which is another way directly, probably the most direct way to give money to a firm. You can only buy a share in a firm. You can go out tomorrow, and you can buy a share in Apple. You get a piece of paper that says you own a share of Apple. You've given Apple your money.

What are they giving you back? They're giving you back two things. The first is they're giving you back the promise that their price of that stock will go up. If you buy-- I don't know, any one-- Apple stock, let's say Apple stock today is selling at 100. I don't know what's selling at. OK. It's probably much more than that, selling at 100.

If I buy it for 100, then what that means I give Apple $100, they give me a share. Why do I do that? Because I think in five years, Apple will be worth $150. And then I could sell that share to someone else and get my $50. In addition, some old style firms, not the new firms, old style firms also sometimes will just send regular checks to people who own their stock, called dividends. That happens less and less.

But basically, so I could put my money in the bank. The bank pays me interest. The firm then borrows from the bank. I can buy a corporate bond. The bond pays me interest. It's a direct loan. Or I can buy stock. In which case I give the money to the firm, and they pay me back by getting more valuable or paying me a dividend. And that's essentially how I transfer money to firms. OK, questions about that? Yeah?

AUDIENCE: Is it really related to these factors, for the interest rate, for example, to the slopes of the demand and supply?

JONATHAN GRUBER: Relate the interest rate to the slope. I don't quite understand.

AUDIENCE: Like you mentioned, the demand curve slopes down because the interest rate goes up.

JONATHAN GRUBER: Yeah, yeah, exactly, there's a mathematical function. Essentially, the demand curve will be a negative function of interest rates. Supply curve will be a positive function of interest rates. There's a mathematical function just like other demand and supply curves. Other questions?

So now, the demand part is uninteresting. We know how to derive this. What you do is, you take a production function. You would get a cost function you could solve for the marginal revenue product of capital. You know how to do that.

The supply side is more interesting. And why is it more interesting? Because once again, substitution income effects are going to mess things up. Just like they did for labor supply, they are going to do it for savings. And they're going to do so in the framework of what we call intertemporal choice, which is your decision of how much do you consume and how much do you save.

So to do that, let's go to figure 16-7. And this diagram describes the person's decision of how much to save. I'm just going to do this graphically. In section, you're going to do the math.

Individuals have a utility function. We talked about the utility function, we choose between two goods. We always do two-element utility function to make the math easy. We have to choose between two goods. We talked about choosing between consumption and leisure.

Now we're talking about choosing between consumption in period one and consumption in period two. Now the trade-off, instead of being between pizza and cookies, or being between work leisure, and consumption, that's what you're consuming today versus consuming tomorrow.

And we understand the trade-off. The trade-off is that basically, people are by and large impatient. If you ask them, would you rather something today or tomorrow, they'd rather have it today. OK, there's the famous marshmallow test they do on little kids, where they see if they'll wait to get a marshmallow. OK, most kids don't. Humans are impatient. OK, we'll come back to that in a few lectures to talk about how you model that. But for now, let's just say humans are impatient.

On the other hand, humans need money in both periods. So let's do a simple example. Imagine my income is $80,000 a year. And imagine I can save that money at 10%. So y is 80,000? And let's imagine I only earn the income-- well, OK, imagine my income is 8,000 a year. I can save that at 10%. And next year, I want to stay home, take care of my kids. I don't want to work.

So what that means is my income in period one will be 80,000, but my income in period two will be zero. Gets more complicated income in period two. You'll see that in the problem set. OK, my income in period two would be zero.

Now, my question is, what do I do? Well, you can see that in figure 16-7. Figure 16-7 has my first period consumption on the x-axis, my second period consumption on the y-axis. One option would be to consume everything in year one and nothing in year two. That's not very attractive.

The other option to consume nothing in year one, in which case, I would have 88,000 in year two. Why? Because I would save 80,000 on a 10% interest. So what's the trade-off? The trade-off is the interest rate, right? Or in other words, what is the price of first period consumption?

The price of first period consumption, is minus-- OK, the price of first year consumption is minus 1 plus i. Once again, I assume all goods have a price of 1. We're not getting into price of goods. This is, what is the relative price of first period consumption?

So assume the price of consumption in period one is 1. The price of consumption in period two is 1. We're normalizing everything to make our math easy. Throwing prices out. OK. It's just we're thinking in terms of goods, effectively, here.

So what is the relative price of consuming in period one? Minus 1 plus i. why is that? Because for every dollar I consume in period one, I give up 1 plus i dollars in period two. That is the opportunity cost of first-period consumption is not having second-period consumption.

Once again, dismal science, every time you eat something, you're giving up the possibility of saving that money, eating more next year. OK, just like the price of leisure is the wage, the price of first-period consumption is minus 1 plus i, or the interest rate, the amount of money you could have the next year.

So basically, that is the trade-off that is representing the temporal choice problem. That's represented by the slope of this budget constraint, which is minus 1 plus i, or minus 1.1 in our example. The slope of the budget rate represents the opportunity cost of first-period consumption in terms of second-period consumption.

Likewise, I have some indifference curve. I have some indifference curve, which represents my preferences over first- and second-period consumption. How impatient I am, basically. So think of that indifference curve basically being how impatient I am.

So like with all of these optimization problems, I know what to do. OK, I find the tangency between the indifference curve, which is essentially my impatience between period one and two, and the budget constraint, which is the market rate at which I can trade-off consumption period one and period two.

Remember, going all the way back to lecture one, the idea of opportunity cost is kind of like the concept of alchemy. I'm not literally turning period one consumption into period two consumption, physically, but I'm effectively doing it through opportunity cost. By not consuming in period one, you get more in period two.

So it's essentially alchemy that happens to the concept of opportunity cost. In this case, it happens through saving in the bank. By saving in the bank, I turn $1 today and turns it into $1.10 tomorrow. $1 in period one into $1.10 in period two. And that's what's happening here.

So my optimal consumption, in this case, is c1 in period one. So what do I do? I say, well, look, given my impatience level, I want to consume c1 in period one in figure 16-7. What that means, I'm going to save 80,000 minus c1. What does that mean my period consumption is? It means my period two consumption is going to be the amount I saved times 1 plus the interest rate. So it's going to be 1 plus i times 80,000 minus c1.

Now, once again, let's remember the trick I've done here without mentioning it. This Is supposed to be about savings. That's the first time I've mentioned savings. Why? Because savings is a bad.

Savings is something you just have to do to get consumption in period two. No one wants to save. This one's even a little more realistic for you guys. Although, you're probably still frugal people. But basically, the idea of savings is not a good thing. We don't have a model where people-- it's not a Scrooge McDuck model where people want to swim in their money. Savings is not a good thing.

Savings is a necessary tool to get consumption from period one to period two. It's a bad. We don't model bads. We model goods. So we don't model savings. We model consumption, and then we get savings as a residual. So when you want to solve for savings, you model consumption, and you get savings of the residual. Just like when you want to model work, you solve for leisure, and you get work as the residual.

Now, let's ask what happens when the interest rate goes up? And once again, I'm sorry. I don't have time for the math here, but we'll do it in section on Friday. Let's ask what happens when the interest rate goes up? Well, we move from bc1 to bc2. That has, guess what, two effects.

On the one hand, you get the substitution effect. The substitution effect is the price of first-period consumption has gone up. So you want to do less of it. Focus on that. Remember, I'm not talking about savings. I'm not talking about second-period consumption.

The substitution income effects are purely in terms of first-period consumption. c1 is what you're maximizing over here. The price of first-period consumption has gone up. Why has it gone up? Because the opportunity cost has gone up. You can have more in period two by saving more. So the substitution effect necessarily says you have less period-one consumption.

But what about the income effect? Well, consumption in period two is a normal good. I'm sorry. Because I'm sorry. Consumption period one is a normal good. Consumption period one is a normal good. And are you richer or poorer if the interest rate goes up? Are you effectively richer or poorer, effectively? Go ahead, you know.

Richer, you're effectively richer because every dollar of savings yields more money. So the income effect says, since you're rich, you want more of everything. What's one of the things you want more of? c21. So the income effect goes the other way. You get the same sort of Giffen good effects without weirdness.

So actually, whether a higher interest rate raised or lowering saving is ambiguous. Once again, think of the intuition. It's even easier intuition here. You're a kid. You want a bike. The bike is going to cost $200. You can earn $10 an hour, so you work 20 hours to get the bike, then you quit. Well, if you're earning 15 bucks an hour, you work fewer hours because all you want is the bike.

That would be a downward-sloping supply curve. In that case, the interest rate rise would cause you-- I'm not working. I'm sorry. Let me back up. I do working. I'm going too fast here.

You want a bike? It's $200. So you got to save $200. You got to save $200. So if the interest rate is 10%, then you have to save, say 180, roughly, this year, to get the bike next year. If the interest rate is 20%, you'd only have to save about 160 this year to get the bike next year. OK forget my other example. I was on labor.

So if you want the bike next year for $200 and interest rate's 10%, you got to save 180. If it's 20%, you got to save 160. So your actual savings goes down. A higher interest rate caused the savings to go down. That would be a downward-sloping supply curve. So just like labor supply, we don't know which way the capital supply curve goes.

And here, unfortunately, unlike labor supply, we don't even have good evidence. We pretty much know that labor-supply curves slope up but are almost vertical. With capital-supply curves, we actually don't know. It's a real empirical weakness in economics.

So we are going to go ahead and assume substitution effects dominate and generally draw upward-sloping capital supply curves in this course. But unfortunately, we don't really know what the capital-supply curve looks like, OK.

I will stop there, and we will come back next time and just continue discussing. We'll discuss how you should invest your money. So come back next time and learn how you should invest your money.